AMD WHITE PAPER

AMD Smarter Choice

Huawei ATCA® Blade Server Powered by AMD Opteron™ Processors

Executive Overview

When AMD first introduced the AMD Opteron™ processor with Direct Connect Architecture in 2003, the target market was for servers in the datacenter. Back then, AMD had a new vision for how system memory and I/O interface to the processor core as well as the ability to create the world's first native dual-core and quad-core x86 processors. These technological advances by AMD have also brought out a new metric called "Performance per watt" which is frequently used as a way to measure overall processor effectiveness. Besides the traditional server market, AMD Opteron processors are also being used for high-end embedded applications. The AMD64 superscalar processor family features Direct Connect Architecture with an integrated memory controller, a scalable high-speed HyperTransport™ technology I/O interface and industry-leading x86 processor performance. The integration of the memory controller and the flexibility of AMD64 technology's Direct Connect Architecture helps facilitate designs with a low chip count, small footprint and low power and cooling requirements at flexible price/performance points. AMD also designs for a stable roadmap from both a thermal and pin-out perspective for generations of processors including quad-core. AMD Opteron processors are NEBS/ETSI friendly and are available in a wide range of power envelopes and with an extended longevity program for select models. These features make the AMD Opteron processor ideal for high-end embedded telecom platforms such as AdvancedTCA® which is currently being deployed for next generation service applications.

This paper describes how AMD Opteron processors have enabled Huawei's ATCA blade server platform to achieve a platform that offers its customers greater flexibility, scalability, performance per watt while helping reduce total operational costs.

AdvancedTCA

Advanced Telecom Computing Architecture (ATCA) is an industry standard specification created by the PCI Computer Manufacturers Group (PICMG) for the next generation of carrier grade communications equipment. ATCA incorporates the latest high speed interconnect technologies, next generation processors and improved reliability, manageability and serviceability features. The ATCA family of standards defines new form factors for full and fractional size system cards, mezzanine cards and a chassis/shelf optimized for communications. ATCA provides standardized platform architecture for carrier-grade telecommunication equipment which can scale in the number of ports, bandwidth, power envelope and physical size based on standard chassis. Additionally, ATCA complies with network operation environmental features such as Network Equipment Building Standard (NEBS), European Telecommunication Standard Institute (ETSI) and meets the 99.999% availability

requirement as demanded by Telecom Equipment Manufacturers. Some reasons why ATCA is ideal for central office applications are because it features high availability, scalability, flexibility and a unified system management approach. For high availability, it incorporates dual redundant components throughout and field replaceable units. Its design allows for both scalability and flexibility because it provides a wide margin of room for compute performance and I/O. It supports multiple fabric architectures and technologies as well as custom interfaces. Some of the end user benefits are that it enables faster time to market, supports a wider range of products, allows for lower barrier to entry or cost efficiencies, and is carrier grade from HW to SW application layers.

Some of the main applications that will be deployed using ATCA based hardware are NGN (Next Generation Networking describes emerging computer network architectures and technologies) and IMS (IP Multimedia Subsystem architectures are those which combine voice/data with additional media such as video). The basic idea behind NGN is that all data and information (voice, data, media, etc.) is transmitted via packets according to service level and security while IMS is an open NGN multimedia architecture for mobile and fixed IP services. Other ATCA applications include Radio Network Controllers for Mobile (RNC), Video/ IPTV and as a subscriber database server (HLR-Home Location Register).

Why AMD Opteron™ Processors for ATCA and High End Telecom Applications?

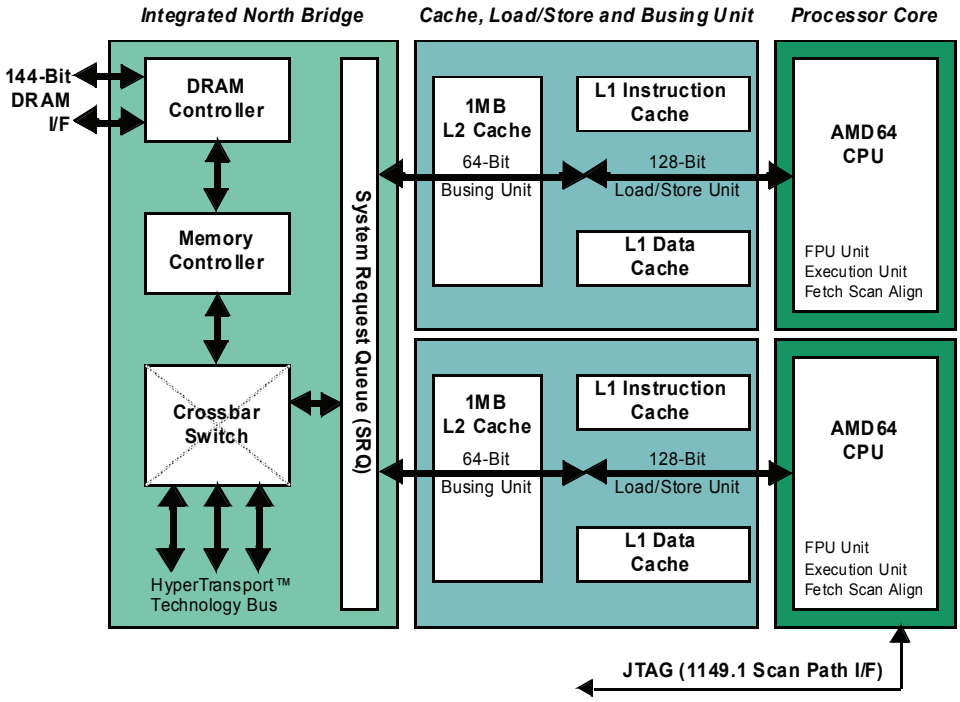

The AMD Opteron processor is well suited for ATCA and high-end telecom applications. From a telecom equipment manufacturer's point of view, total system performance is much more than CPU clock speed. They measure performance based on how much I/O can be streamed without degrading memory performance, how much work can be done in "n" clock cycles at "x" watts as well as how much board real estate is needed to meet the needs of the application. These factors along with other telecommunications industry and ATCA-specific "care abouts" fit well with the entire AMD Opteron processor and ecosystem value proposition. These include:

- **Total Cost of Ownership** – As pointed out above, one of the main benefits of ATCA to end users is total cost of ownership. Not just the price of the processor, but the total

system cost system, including the cost to operate and maintain it. Another measurement of cost is price per channel which relates closely to I/O bandwidth of the system, which the AMD Opteron processor and HyperTransport™ technology address exceptionally well.

- **Reduced Risk** – The ability to demonstrate that an AMD Opteron processor will meet the needs of the mission (i.e., the high availability of platforms).
- **Performance per Watt** – The measure of system performance when compared with the average system power usage. This is done by taking the performance score from a benchmark or application and dividing it by the average system power usage. Performance per watt is quickly becoming one of the most integral design paradigms for x86 processors for all markets. The AMD Opteron processor and HyperTransport technology address this measurement exceptionally well. In fact, all three AMD Opteron processor thermal design power (TDP) envelopes fall within the ATCA specification.
- **Stability and Investment Protection** – Telecom equipment providers design systems that require longevity support and a stable platform. Select models of the AMD Opteron processor family have longevity support for up to seven (7) years (5+2). The current AMD Opteron processor family based on DDR2 memory has been designed for stability from both a thermal and pin-out perspective to be compatible over several generations of multi-core processors. Hence, a platform engineer can design a system today based on a 90nm dual-core Second-Generation AMD Opteron processor, which is designed to be upgradeable to the next generation true quad-core technology.
- **Scalability** – AMD64 with Direct Connect Architecture allows the designer to scale the system by adding multiple processors and/or co-processors through HyperTransport technology connected directly to other AMD Opteron processors in a modular fashion. Each additional processor can add incremental I/O and memory bandwidth.
- **OS Agnostic** – The ability to support multiple CGOS images (i.e., Solaris, UNIX, Windows®, Linux®, etc.).
- **Carrier Grade Reliability, Availability and Serviceability or Carrier Grade RAS** – AMD's Direct Connect Architecture can provide the necessary RAS features needed for next generation telecom equipment including ECC protected on-chip and off-chip memory arrays, memory sparing, chip kill, CRC for in flight data and an inherent fault tolerant architecture with excellent MTBF.
- **Standard Compliance such as NEBS and ETSI** – The ability to meet or exceed the Network Equipment Building System (NEBS) requirements or ETSI (European Telecom. Standard Institute the NEBS equivalent for Europe). This means the ability to meet the strict thermal requirements of NEBS (i.e., to continue to run at full clock rate in a 55DEG-C ambient environment for up to 96-hours).
- **Support** – AMD is committed to our Telecom customer's success and offers reliable design and manufacturing, longevity of components and design support.

About Huawei

Huawei Technologies is a leader in providing next generation telecommunications networks. The company is committed to providing innovative and customized products, services and solutions to create long-term value and potential growth for its customers. Huawei now serves 31 of the world's top 50 operators, along with over one billion users worldwide. Huawei's products and solutions encompass wireless, core network, network equipment, system level applications / software as well as terminals.

Huawei's global contract sales for 2006 reached $11 billion (USD), 65% of which comes from the international market. 2006 global contract sales represent a year-over-year growth rate of over 34%. Huawei has now become a leading vendor in the industry. As far as market share is concerned, Huawei is No. 2 in worldwide share of optical network hardware market; No. 2 in worldwide share of broadband access market (DSL Port Shipments); No. 3 in worldwide share of the carrier Ethernet switch/router market. Huawei has also become one of the few vendors in the world to provide end-to-end 3G solutions.

At the end of 2006, Huawei has over 61,000 employees, of whom 48% are dedicated to R&D. Huawei's global R&D centers are located in Bangalore in India, Silicon Valley and Dallas, TX in the US, Stockholm in Sweden and Moscow in Russia – this is in addition to those in Beijing, Shanghai, Nanjing, Shenzhen, Hangzhou and Chengdu China.

Huawei Tecal T8000 Computing Platform

The Huawei Tecal T8000 platform is an ATCA chassis based computing platform that includes full hardware, operating system, middleware, system management and application layer integration. It is designed for telecommunications and enterprise/large scale networks that need high availability and it offers the flexibility to allow for customization and scalability. The Tecal 8000 helps customers to consolidate their resources through a centralized management system. This leads to lower operational costs and maximizes overall efficiencies. One of the main advantages of the Tecal 8000 is its ability to save energy through its low power consumption. A large portion of this power savings is due to the ability of the system with its CPU processors to automatically adjust the processors' operating frequency according to workload usage. The T8000 also uses a centralized power supply system and effective AC-DC power converter modules to gain nearly 90% conversion efficiency. The cooling and internal temperature regulation in the T8000 is such that it can automatically be controlled through the management module system. By using AMD Opteron processors (which have inherent telecom RAS and ECC capabilities) as well as other core components with redundant design, hot swappable support, and double management backup, the Tecal T8000 offers high reliability and availability (99.999%). Other reliability features include automatic power control protection to prevent power overload and electronic error interpretation functions.

The Tecal T8000 platform is a modular design that enables the sharing of resources and allows users the ability to have a stable and flexible on-demand configuration. AMD provides Huawei and their customers with additional investment protection by offering 5+2 years longevity support on select processors and a stable socket infrastructure allowing for seamless upgrades to new generations of AMD Opteron processors. The T8000 can be customized easily for different computing, memory performance and external data storage and network adaptability through the reconfiguration of a small number of components. For example, using the Huawei BA22 ATCA server board, customers can upgrade their T8000 platform for greater performance with the replacement of a dual-core AMD Opteron processor control blade with a true quad-core processor seamlessly. With the BA22 server board, customers can upgrade from a dual socket AMD Opteron processor configuration to a 4 socket platform resulting in increased I/O and memory bandwidth. The T8000 features the ability to support Gigabit Ethernet, Fibre Channel, SAS, SCSI and other interfaces through the extended interface module. There is support for multiple versions of Linux including Carrier-grade Linux support (CGL2.0), Windows Operating System and full compatibility between 64-bit and 32-bit applications. Huawei provides field and service support that helps with the deployment process, software distribution and the management system. The management Software supports equipment management, fault diagnosis, flow and load balancing, firmware upgrade, alarm monitoring, asset management etc. The T8000 chassis can support up to 14 vertically placed ATCA blades. Typical configurations can include blades that provide switching, control (server blades), load-balancing, storage and network security intrusion.

Huawei has selected AMD64 solutions as its processor of choice for its T8000 compute platform server blade because of its overall outstanding performance and the support provided by AMD engineering. The Huawei BA40 server blade uses four 800 series dual-core AMD Opteron processors along with eight memory DIMMs - offering 16GB of DDR PC3200 SDRAM (running at 400MHz). It has support for two 2.5 inch hot swappable SAS Hard Disk Drives and a RAID1. The Huawei BA22 server blade uses two AMD Opteron 2000 Series processors or upcoming quad-core processors supported by four DIMMs for 16GB of DDR2 SDRAM. It supports GigE, SAS, Fibre Channel for intra chassis communication. For external I/O it supports two USB 2.0, RS232, up to eight GigE, two Fibre channel and two SCSI interfaces.

BA22 Specifications

Processor	Second-Generation AMD Opteron™ processor Socket F (1207)
Quantity of Processors	1/2 (standard configuration/maximum configuration)
Cache (Maximum)	1MB L2/inner core
Memory	Up to 16 GB DDR2 SDRAM. Support Register, ECC, SDCC (Chip-Kill) and memory mirror. The working frequency is 667 MHz.
Built-in Disk Driver	None. (Expand two 2.5-inch SAS/SATA disks through optional server boards to support RAID 0/1 and hot swap.)
Intra-chassis Interconnection Interface	4* GE, 2*FC, 2*SAS
External I/O	Two USB 2.0 interfaces, one RS232 serial port. Able to expand various interface modules, such as GE, FC and SCSI as well as the local KVM interface. The maximum number of interfaces can be expanded to 8*GE/6*GE+2*FC/6*GE+2*SCSI.
System Management	Integrates the BMC management module. BMC connects to the SMM through IPMB, and provides the management functions, such as monitoring, alarm, log and asset.
Supported Operating Systems	Windows® Server 2003 Enterprise Edition SP1 Windows Server 2003 Enterprise x64 Edition SP1 SuSE Linux Enterprise Server 9 SP2 for x86/AMD64/Intel EM64T SuSE Linux Enterprise Server 9 SP3 for x86/AMD64/Intel EM64T SuSE Linux Enterprise Server 10 for x86/AMD64/Intel EM64T Red Hat Enterprise Linux AS 3 Update 5 for x86/AMD64/Intel EM64T Red Hat Enterprise Linux AS 3 Update 7 for x86/AMD64/Intel EM64T Red Hat Enterprise Linux AS 3 Update 8 for x86/AMD64/Intel EM64T Red Hat Enterprise Linux AS 4 Update 2 for x86/AMD64/Intel EM64T Red Hat Enterprise Linux AS 4 Update 3 for x86/AMD64/Intel EM64T Red Hat Enterprise Linux AS 4 Update 4 for x86/AMD64/Intel EM64T
Supported Database	Oracle, Informix, SQL2000, DB2, etc.
Standards	Conforms to ATCA®, RoHS, CE and UL standards.

BA22 Component

Item		Description
Server Interface Board	BR25*	Dual disk storage, support expansion of one interface module.
	BR26*	No storage. Support expansion of four interface modules. In which, class I supports up to 1.
	BR28	Flash storage. Does not support expansion interface module.
Server Interface Board Interface Module	Class I	2G FC interface module, connect to 2*FC externally SCSI interface module, connect to 2*SCSI Ultra320 externally Fabric 2G FC interface module (connect to backplane internally)
	Class II	GE interface module, connect to 2*GE externally Fabric GE interface module (connect to backplane internally)

*Note: BR25 does not support Fabric GE interface module. BR26 does not support repeated configuration of interface modules of any internally connected backplane, other interface modules can be configured repeatedly.

BA40 Component

Processor	Support up to four Dual-Core AMD Opteron™ 800 Series processors using 55W low power consumption
Memory	Standard configuration: Eight DIMM slots, supporting up to 16GB Register DDR SDRAM and supporting ECC and SDCC. The working frequency is 400MHz.
Internal Storage	Supports two 2.5-inch SAS hard disk, hot swap and RAID1.
Port	Standard configuration: two 10/100/1000 MB Ethernet port to backplane (Base) Standard configuration: two 1000 MB Ethernet ports to backplane (Fabric) Provides two 10/100/1000 MB Ethernet ports to peripheral equipment (optional, provided through the server interface board)
Fiber Channel (Optional)	Two 1 G/2G FC interfaces, support FC-AL and FC-SW.
Internal Expansion (optional)	Fabric GE&FC pinch plane, integrating the GE and FC controller
VGA	Integrated video controller
I/O Port	Standard configuration: two USB 2.0 interfaces Standard configuration: one RS232 interface
Equipment Management	Conform to IPMI 1.5 specifications. Support board hot swap, monitoring, alarm, log and asset management functions.
Equipment Indicator	OOS, HEALTHY, HDD, HOTSWAP, etc.
Supported Operating Systems	Windows Sever 2003 Enterprise Edition SP1 Windows Sever 2003 Enterprise x64 Edition SuSE Linux Enterprise Server 9 SP3 for x86/AMD64/Intel EM64T Red Hat Enterprise Linux AS 4 Update 3 for x86/AMD64/Intel EM64T Red Hat Enterprise Linux AS 3 Update 5 for x86/AMD64/Intel EM64T Red Hat Enterprise Linux AS 3 Update 7 for x86/AMD64/Intel EM64T
Supported Database	Oracle, Informix, SQL2000, DB2

Conclusion

In the past, telecommunications equipment manufacturers and network equipment providers often relied on increased CPU clock speed, bandwidth, and/or storage capacity for solving compute issues on the network. However, with today's new telecom business landscape – a converged industry with fewer, more competitive players; constrained resources and revenue; growing demand in customer numbers and services demanded; the rise of IP-based applications, and the challenges of building out and upgrading infrastructure, including power – TEMs and NEPs must focus energies on system efficiency and finding ways to increase data throughput within the system – a problem that can be addressed by innovations such as HyperTransport technology. Huawei's ATCA Tecal 8000 blade server powered by the AMD Opteron processor offers a platform solution with flexibility, high availability, centralized management, scalability and future investment protection.

For more information on the embedded AMD64 products and programs, please visit www.amd.com/amd64embedded.

For more information on Huawei Technologies, please visit www.huawei.com

AMD Smarter Choice

Advanced Micro Devices
One AMD Place
P.O. Box 3453
Sunnyvale, CA 94088-3453
www.amd.com

43319-A